arXiv:1604.04473v1 [cs.CV] 15 Apr 2016

Probing the Intra-Component Correlations within Fisher Vector for Material Classification[☆]

Xiaopeng Hong[a,*], Xianbiao Qi[a,b], Guoying Zhao[a], Matti Pietikäinen[a]

[a] *Department of Computer Science and Engineering, University of Oulu, Finland*
[b] *Beijing University of Posts and Telecommunications*

Abstract

Fisher vector (FV) has become a popular image representation. One notable underlying assumption of the FV framework is that local descriptors are well decorrelated within each cluster so that the covariance matrix for each Gaussian can be simplified to be diagonal. Though the FV usually relies on the Principal Component Analysis (PCA) to decorrelate local features, the PCA is applied to the entire training data and hence it only diagonalizes the *universal* covariance matrix, rather than those w.r.t. the local components. As a result, the *local decorrelation* assumption is usually not supported in practice.

To relax this assumption, this paper proposes a completed model of the Fisher vector, which is termed as the Completed Fisher vector (CFV). The CFV is a more general framework of the FV, since it encodes not only the variances but also the correlations of the whitened local descriptors. The CFV thus leads to improved discriminative power. We take the task of material categorization as an example and experimentally show that: 1) the CFV outperforms the FV under all parameter settings; 2) the CFV is robust to the changes of the number of components in the mixture; 3) even with a relatively small visual vocabulary the CFV still works well on two challenging datasets.

Keywords: Image representation, Fisher vector, Gaussian mixture model,

[☆]This work is supported by the Academy of Finland and Infotech Oulu

[*]Corresponding author
Email addresses: xiaopeng.hong@ee.oulu.fi (Xiaopeng Hong), qixianbiao@gmail.com (Xianbiao Qi), gyzhao@ee.oulu.fi (Guoying Zhao), mkp@ee.oulu.fi (Matti Pietikäinen)

1. Introduction

Image representations, which extract informative cues from images, play a central role in various computer vision tasks such as material classification and object recognition. With the increasing demand for 'ideal' image representations which are both highly discriminative in distinguishing instances from classes and well robust to the possibly large intra-class divergence in appearance, a number of image representations have been developed in the literature [1, 2].

The bag-of-visual words (BoW) model has become a popular framework for image representations. The BoW extracts the local features, which characterize the *local* properties of images, and assigns them to the closed entries in a codebook. A codebook, more commonly known as a vocabulary [3, 4, 5, 6, 7, 8, 9], is usually generated by descriptor quantization, such as automatic vector quantization [4, 5, 6, 7, 10] and manual design [3, 8]. The BoW uses the histogram of visual words as the image representation, to approximate the occurrence frequency of visual words (in the codebook) for a particular image. In general, the BoW is robust against noise and background clutters. However, in order to make the vocabulary compact, the descriptor quantization usually learns or selects a rather small codebook (typically in a size of 200 to 1,000) to represent hundreds of thousands of local descriptors. It thus inevitably leads to considerable information loss [11, 12].

Recently, the BoW framework has been extended to the approaches which aggregate the local descriptors into a vector representation incorporating additional information [13, 14, 15, 16, 17]. The Fisher vector (FV) [2, 11, 13, 18, 19, 20, 21, 22], as one successful extension of the BoW, has aroused a great deal of research interest. The FV models the visual words via a Gaussian mixture model (GMM). Given an image represented by a group of local descriptors, the FV measures the deviations of local descriptors (from this image) w.r.t. the GMM parameters. Compared with the BoW, the FV brings additional infor-

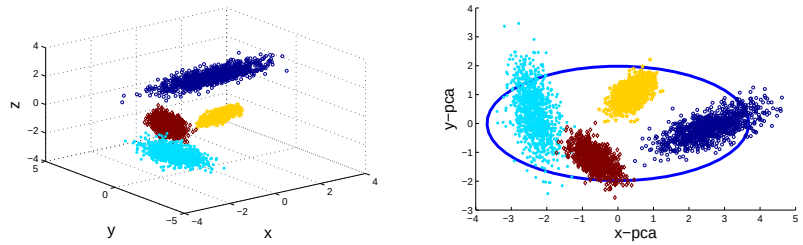

(a) 3D sample points from four-component GMM

(b) Projection to 2D sub-space by PCA

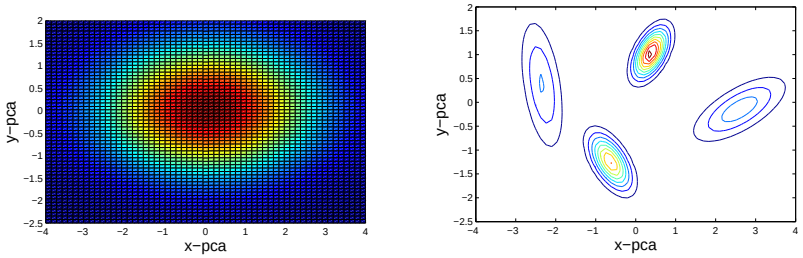

(c) The *universal* covariance matrix is diagonal

(d) Features are correlated within each component

Figure 1: Universal PCA fails to eliminate the correlation within local modes

mation, thus leads to substantial improvement in accuracy. Moreover, usually a smaller vocabulary is required for the FV to achieve satisfying accuracy. In addition, the rich dimension of the FV makes the corresponding image representations more linearly separable. Hence the FV performs well with the efficient linear classifiers. The superiority of the FV has been shown in [1, 2, 23], where state-of-the-art feature aggregating approaches are empirically evaluated and the Fisher vector is suggested to be the most competitive one. Another example can also be found in [24], in which the Fisher vector based systems obtain the best results in the FGCOMP fine-grained Challenge 2013.

Despite of its great success, the existing Fisher vector framework is based on an underlying assumption that local descriptors are well decorrelated within each cluster (*Local Decorrelation Assumption*). All covariance matrices of the components in the mixture are correspondingly simplified and forcedly restricted to be diagonal. Accompanied by several approximations, the assumption signifi-

cantly eases the formulation of the FV [18, 20]. However, as a result, it severely restricts the FV to accurately characterizing the local features for images in complicated situations.

One might argue that the Fisher vector usually relies on the well-established Principal component analysis (PCA) to decorrelate local features and also reduce the dimensionality [20]. However, it should not be forgotten that the PCA is applied to the entire training samples. Thereby it is the *universal* covariance matrix (of the entire training samples) that the PCA only diagonalizes. There is no guarantee of the decorrelation for the local features from training images within each GMM component. Nor is it for those from testing ones.

Figure 1 provides an example to illuminate the limitations of the *Local Decorrelation Assumption* in the FV. Figure 1(a) shows points sampled from a four-component GMM in the three-dimensional space. In Figure 1(b), they are projected to the 2D subspace by the PCA. Note that the PCA successfully decorrelates the entire set of samples and as a result diagonalizes the *universal* covariance matrix. It is demonstrated in Figure 1(b) by the *non-sloping* blue ellipse, which stands for the contour line of the Gaussian probability density function on all points. Consistent observations can also be made from Figure 1(c), where the plotted surface of the Gaussian probability density with the *universal* mean and covariance matrix is not sloping either. However, the decorrelation assumption is not supported when taking each Gaussian component into account separately. In Figure 1(d), the apparent skewness of the contour line of each local Gaussian serves as a reminder of the undisputed existence of the feature correlations within each cluster. Obviously the *local decorrelation* assumption is strictly satisfied in rare cases when the orientations of the principle axes for all GMM components happen to be close to those of the uni-modal Gaussian on all points. However in much more common cases it is violated if the orientations of the principle axes vary, as shown by the example here. A practical example in material classification is also provided to further illustrate this problem in Section 3.1.

Urged by the observation we reached, in this paper we propose to explore

the feature correlations for improving the FV. We revisit the Fisher vector from the perspective of whitening transformation, and formally derive a vector representation taking into account feature correlations within the GMM clusters rather than only the variances (as the FV does). The resulting representation is named the Completed Fisher vector (CFV). It is a more general case of the FV where the *local decorrelation* assumption is relaxed. In the following sections, we evaluate the proposed CFV and compare it with its counterpart namely the FV on two challenging data set including KTH-TIPS2 and FMD. Superior performance on these two data sets successfully validate the effectiveness of the proposed CFV.

The remaining of this paper is organized as follows: In Section 2, we derive our completed Fisher vector through in depth revisiting the standard Fisher vector. In Section 3.1, we provide a comparative study between the completed FV and the traditional diagonal ones on two common used local descriptors under different dimensions. The comparisons with the state of the art methods on two publicly available datasets are reported in Section 3. Finally, Section 4 concludes the whole paper.

2. Exploring Intra-Component Correlations

In this section, the existing Fisher vector is firstly introduced to facilitate the derivation of the completed Fisher vector in Section 2.1. In Section 2.2 we revisit the FV from the viewpoint of whitening transformation. Based on this understanding, we formulate the CFV and discuss on the implementation issues in Section 2.3.

2.1. Standard Fisher Vector

The Fisher vector [18] usually generates the vocabularies of visual words by means of the GMM. Let the set of parameters of a K-mode GMM be $\Theta = \{\boldsymbol{\pi}, \boldsymbol{\mu}, \boldsymbol{\Sigma}\}$, where $\boldsymbol{\pi} \equiv \{\pi_1, \cdots, \pi_K\}$, $\boldsymbol{\mu} \equiv \{\boldsymbol{\mu}_1, \cdots, \boldsymbol{\mu}_K\}$, and $\boldsymbol{\Sigma} \equiv \{\boldsymbol{\Sigma}_1, \cdots, \boldsymbol{\Sigma}_K\}$ are the sets of the priors, mean vectors, and covariance matrices of the K components respectively. Denote by $I = \{\mathbf{x}_n, n = 1 \cdots N\}$ the set of D-dimensional

feature vectors extracted from an image. The GMM then associates each $\mathbf{x}_n$ with the *k-th* mode in the mixture via a soft assignment given by the posterior probability:

$$\lambda_{nk} = \frac{\pi_k \mathcal{N}(\mathbf{x}_n, \boldsymbol{\mu}_k, \boldsymbol{\Sigma}_k)}{\sum_{t=1}^{K} \pi_t \mathcal{N}(\mathbf{x}_n, \boldsymbol{\mu}_t, \boldsymbol{\Sigma}_t)}, k = 1, \cdots, K, \tag{1}$$

where $\mathcal{N}(\mathbf{x}, \boldsymbol{\mu}_k, \boldsymbol{\Sigma}_k)$ for $k = 1, \cdots, K$ are the component Gaussian density functions of $\mathbf{x}$ with mean vector $\boldsymbol{\mu}_k$ and covariance matrix $\boldsymbol{\Sigma}_k$.

Under the *local decorrelation* assumption, the covariance matrices of components are simplified to be diagonal, i.e. $\boldsymbol{\Sigma}_k \overset{\text{s.t.}}{=} \boldsymbol{\Lambda}_k \equiv \text{diag}\boldsymbol{\sigma}_k^2$, where $\boldsymbol{\sigma}_k \in \mathbb{R}_+^D$. The Fisher vector of image I is denoted by $\mathcal{F}(I) = \left[\cdots \mathbf{u}_k^T \cdots \mathbf{v}_k^T \cdots\right]^T$, where $\mathbf{u}_k$ and $\mathbf{v}_k$ represent the first-order and second-order statistics of the deviations with respect to the parameters of the *k-th* Gaussian [13, 18, 20]. More specifically, the d-th entries of $\mathbf{u}_k$ and $\mathbf{v}_k$ for $d = 1, \cdots, D$ are defined as follows:

$$u_{dk} = \frac{1}{N\sqrt{\pi_k}} \sum_{n=1}^{N} \lambda_{nk} \frac{x_{dn} - \mu_{dk}}{\sigma_{dk}}, \tag{2}$$

$$v_{dk} = \frac{1}{N\sqrt{2\pi_k}} \sum_{n=1}^{N} \lambda_{nk} \left[\left(\frac{x_{dn} - \mu_{dk}}{\sigma_{dk}} \right)^2 - 1 \right]. \tag{3}$$

Let $\lambda_k = \sum_{n=1}^{N} \lambda_{nk}$ and re-parametrize the soft-assignment as $\bar{\lambda}_{nk} = \frac{\lambda_{nk}}{\lambda_k}$. Because $\boldsymbol{\Sigma}_k \overset{\text{s.t.}}{=} \boldsymbol{\Lambda}_k$, Eqs. 2 – 3 can be expressed by the following multivariate form:

$$\mathbf{u}_k \propto \lambda_k \sum_{n=1}^{N} \bar{\lambda}_{nk} \left[\boldsymbol{\Lambda}_k^{-\frac{1}{2}} (\mathbf{x}_n - \boldsymbol{\mu}_k) \right], \tag{4}$$

$$\mathbf{v}_k \propto \lambda_k \sum_{n=1}^{N} \bar{\lambda}_{nk} \left\{ \left[\boldsymbol{\Lambda}_k^{-\frac{1}{2}} (\mathbf{x}_n - \boldsymbol{\mu}_k) \right] \circ \left[\boldsymbol{\Lambda}_k^{-\frac{1}{2}} (\mathbf{x}_n - \boldsymbol{\mu}_k) \right] - \mathbf{1} \right\}, \tag{5}$$

where $\mathbf{1}$ is the unit vector and the binary operator $\circ$ denotes the Hadamard product, a.k.a the entry-wise product.

The formal but involved derivations of the Fisher vector are provided in [13, 20] from the viewpoint of the normalised gradients of the Fisher information matrix. We notice even under the *local decorrelation* assumption, Eqs. 2 – 5 can only be obtained by approximating the Fisher information matrix and the

partial derivatives w.r.t. its parameters. In attempting to derive a more general form of the FV by incorporating the intra-component feature correlations, it would be inevitable to make more assumptions and approximations if along the same line of thought. To avoid being enmeshed in such difficulties, we provide a second interpretation for the FV from the perspective of whitening transformation [25, 26], which provides a natural and more efficient way to introduce the corresponding formulas.

2.2. Fisher Vector: A Revisit

In probability theory and statistics, consider the D-dimensional vectors in a set $I = \{\mathbf{x}_n, n = 1 \cdots N\}$ as discrete random variables. We have the expected value $\hat{\boldsymbol{\mu}}$ and covariance matrix $\hat{\boldsymbol{\Sigma}}$ as follows:

$$\hat{\boldsymbol{\mu}} = E\left(\mathbf{x}\right) = \sum_{n=1}^{N} p\left(\mathbf{x}_n\right) \mathbf{x}_n, \tag{6}$$

$$\hat{\boldsymbol{\Sigma}} = E\left(\left(\mathbf{x} - E\left(\mathbf{x}\right)\right)\left(\mathbf{x} - E\left(\mathbf{x}\right)\right)^T\right) = \sum_{n=1}^{N} p\left(\mathbf{x}_n\right)\left(\mathbf{x}_n - \hat{\boldsymbol{\mu}}\right)\left(\mathbf{x}_n - \hat{\boldsymbol{\mu}}\right)^T, \tag{7}$$

where $p\left(\mathbf{x}\right)$ is an underlying probability mass function (pmf) of $\mathbf{x}$, s.t. $p\left(\mathbf{x}\right) \geq 0$ and $\sum_{n=1}^{N} p\left(\mathbf{x}_n\right) = 1$ (*the 'Completeness' property*).

Without losing generality, suppose $\hat{\boldsymbol{\Sigma}}$ to be positive definite. By left multiplying the inverse square root of $\hat{\boldsymbol{\Sigma}}$ on both sides of Eq. 6, and according to the distributivity of matrix multiplication we have

$$\sum_{n=1}^{N} p\left(\mathbf{x}_n\right) \hat{\boldsymbol{\Sigma}}^{-\frac{1}{2}} \mathbf{x}_n - \hat{\boldsymbol{\Sigma}}^{-\frac{1}{2}} \hat{\boldsymbol{\mu}} = \sum_{n=1}^{N} p\left(\mathbf{x}_n\right) \hat{\boldsymbol{\Sigma}}^{-\frac{1}{2}} \mathbf{x}_n - \sum_{n=1}^{N} p\left(\mathbf{x}_n\right) \hat{\boldsymbol{\Sigma}}^{-\frac{1}{2}} \hat{\boldsymbol{\mu}} \tag{8}$$

$$= \sum_{n=1}^{N} p\left(\mathbf{x}_n\right) \underbrace{\left[\hat{\boldsymbol{\Sigma}}^{-\frac{1}{2}}\left(\mathbf{x}_n - \hat{\boldsymbol{\mu}}\right)\right]}_{\mathbf{w}_n} = \mathbf{0} \quad , \tag{9}$$

Similarly, by left and right multiplying $\hat{\boldsymbol{\Sigma}}^{-\frac{1}{2}}$ on both sides of Eq. 7, the equation becomes:

$$\sum_{n=1}^{N} p\left(\mathbf{x}_n\right) \hat{\boldsymbol{\Sigma}}^{-\frac{1}{2}}\left(\mathbf{x}_n - \hat{\boldsymbol{\mu}}\right)\left(\mathbf{x}_n - \hat{\boldsymbol{\mu}}\right)^T \hat{\boldsymbol{\Sigma}}^{-\frac{1}{2}} = \mathbf{I}. \tag{10}$$

Recall that $\hat{\boldsymbol{\Sigma}}^{-\frac{1}{2}}$ is symmetric such that $\hat{\boldsymbol{\Sigma}}^{-\frac{1}{2}} = [\hat{\boldsymbol{\Sigma}}^{-\frac{1}{2}}]^T$, Eq. 10 is equivalent to

$$\sum_{n=1}^{N} p\left(\mathbf{x}_n\right) \hat{\boldsymbol{\Sigma}}^{-\frac{1}{2}} \left(\mathbf{x}_n - \hat{\boldsymbol{\mu}}\right) \left(\hat{\boldsymbol{\Sigma}}^{-\frac{1}{2}} \left(\mathbf{x}_n - \hat{\boldsymbol{\mu}}\right)\right)^T = \mathbf{I}. \tag{11}$$

Taking into account the *Completeness* constraint of *pmf*, we rewrite Eq. 11 by

$$\sum_{n=1}^{N} p\left(\mathbf{x}_n\right) \left[\hat{\boldsymbol{\Sigma}}^{-\frac{1}{2}} \left(\mathbf{x}_n - \hat{\boldsymbol{\mu}}\right) \left(\hat{\boldsymbol{\Sigma}}^{-\frac{1}{2}} \left(\mathbf{x}_n - \hat{\boldsymbol{\mu}}\right)\right)^T - \mathbf{I}\right] = \mathbf{0}. \tag{12}$$

By comparing Eqs. 9 and 12 with Eqs. 4 and 5 respectively, we reach the following remarks:

1) It is well known that a linear transformation on the feature space will convert an arbitrary normal distribution into another normal distribution [25]. As revealed by Eqs. 9 and 11, an *ideal* whitening process utilizing the expected value and the covariance matrix of $\{\mathbf{x}_n\}$, translates, scales, and rotates the axes so that the resulting multivariate variable $\mathbf{w}_n = \hat{\boldsymbol{\Sigma}}^{-\frac{1}{2}} \left(\mathbf{x}_n - \hat{\boldsymbol{\mu}}\right) \sim \mathcal{N}\left(\mathbf{0}, \mathbf{I}\right)$, i.e. the standard Gaussian distribution. However, in reality, for example in a GMM, the parameters are optimized on the local features from the entire training set rather than from a single image. It is thus difficult or even impossible for such pre-determined parameter settings to perfectly fit the expected mean and covariance matrix for a particular image I. As a result, the zero vector and the zero matrix on the right sides of Eqs. 9 and 12 usually cannot be reached after whitening transformation.

2) The Fisher vector reflects the remainders as indicated by Eqs. 4 and 5. It measures the misfit between the parameter settings of each Gaussian component and the set of local features from a particular image by aggregating the first and second order statistics of the differences between the whitened vectors, and the zero vector as well as the identity matrix. It is rational that the whitening transformation prevents certain features from dominating differences calculations merely because they have large numerical values.

3) Apparently, the whitening transformation in the FV is a special case of the general one as it uses only diagonal covariance matrices. It simplifies the calculation, yet restricts the form of the probability density function and greatly

limits the power to capture the important feature correlations from data. To remove these limitations, we propose to explore the intra-component correlations in the next section.

2.3. Exploring Intra-Component Correlations

In this subsection we remove the *local decorrelation* assumption and relax the constraint that the covariance matrix should be diagonal. In the light of the relation between Eqs. 9 and 12 and Eqs. 4 and 5, we define the derivations considering the feature correlations inside the GMM component as follows:

$$\tilde{\mathbf{u}}_k = \frac{1}{N\sqrt{\pi_k}} \sum_{n=1}^{N} \lambda_{nk} \left[\boldsymbol{\Sigma}_k^{-\frac{1}{2}} \left(\mathbf{x}_n - \boldsymbol{\mu}_k \right) \right], \tag{13}$$

$$\tilde{\mathbf{V}}_k = \frac{1}{N\sqrt{\pi_k}} \sum_{n=1}^{N} \lambda_{nk} \left[\boldsymbol{\Sigma}_k^{-\frac{1}{2}} \left(\mathbf{x}_n - \boldsymbol{\mu}_k \right) \left(\boldsymbol{\Sigma}_k^{-\frac{1}{2}} \left(\mathbf{x}_n - \boldsymbol{\mu}_k \right) \right)^T - \mathbf{I} \right]. \tag{14}$$

To form a vector representation, we unpack the entries in the matrix $\tilde{\mathbf{V}}_k$ into a vector $\tilde{\mathbf{v}}_k$. The number of entries in $\tilde{\mathbf{v}}_k$ can be reduced as the right side of Eq. 14 is symmetric. In specific, we have $\tilde{\mathbf{v}}_k = \left[\mathbf{vd}_k^T, \alpha \cdot \mathbf{vt}_k^T \right]^T$, where $\mathbf{vd}_k$ is a $D \times 1$ vector consisting of all diagonal entries of $\tilde{\mathbf{V}}_k$, $\mathbf{vt}_k$ includes the $D\left(D-1\right)/2$ elements above the diagonal of $\tilde{\mathbf{V}}_k$, and α is a parameter which balances the contributions of $\mathbf{vd}_k$ and $\mathbf{vt}_k$. For each of the K modes in the Gaussian mixtures, the whitened deviations to the standard normal distribution are completely specified by the $D + D\left(D+1\right)/2 = D\left(D+3\right)/2$ parameters including the elements of the first order mean vector and the full second order covariance matrix.

As a result, the vector representation of image I is the stacking of the vectors $\tilde{\mathbf{u}}_k$ and then of $\tilde{\mathbf{v}}_k$ for each of the K modes in the Gaussian mixtures, i.e., $\tilde{\mathcal{F}}(I) = \left[\cdots \tilde{\mathbf{u}}_k^T \cdots \tilde{\mathbf{v}}_k^T \cdots \right]^T$.

In contrast with the FV, there is no assumption in the form of the component covariance matrix in the CFV. The second order part $\tilde{\mathbf{V}}_k$ is a $D \times D$ matrix covering both the variances and the correlations, while the one of the FV is a $D \times 1$ vector including only the variances, namely the entries in the diagonal of $\tilde{\mathbf{V}}_k$. The resulting vector representation $\tilde{\mathcal{F}}$ is therefore a completed model of

the Fisher vector by the second statistics. It is named the Completed Fisher vector (CFV) in this paper. Obviously, the FV is a special case of the CFV. Only when the diagonal entries of $\tilde{\mathbf{V}}_k$ are used, the CFV degenerates to the FV.

Compared with the Fisher vector whose dimension is $2DK$, and the BoW whose dimension is only K, the CFV has a richer dimension of $D(D+3)K/2$. In general then, knowledge of the covariance matrix allows to reflect the dispersion of the data in any direction, or in any subspace. It thus contains significantly more information from the intra-component correlations. In addition, it can also be understood as an embedding of the local features in a much higher-dimensional space which is more separable for linear classification.

In practice, before using the representation in a linear model (e.g. a support vector machine), we follow the two issues by the improved Fisher Vector (IFV) [18, 20], including: 1) Power normalization, which applies the function $|z|^{\gamma}\operatorname{sign}(z)$ to each entry of the vector, where $0 \leqslant \gamma \leqslant 1$. It can be regarded as a post-processing procedure using Gamma correction. More in-depth study can be found in [27, 28]. 2) L2 Normalization, which implicitly discards the background information as suggested by [18, 20].

To our best knowledge, there are two relevant works which more or less realize the limitations of the aforementioned *Local Decorrelation Assumption* [16, 17]. Picard and Gosselin remove the GMM procedure by using the k-means clustering directly, and straightly use the second second order information to characterize the clusters [16]. In [17], Delhumeau et al. learn a local PCA rotation matrix for each cluster of the partitioned feature space. However, both two works are applied to the VLAD (namely the *vector of locally aggregated descriptors*) framework rather than the FV, which is usually shown with better performance [2]. As a result, for example, no whitening issue is discussed by their works, while in this paper, we derive our improved descriptor from the perspective of whitening. More details in the significant differences between the FV and the VLAD can be found in related publications [2, 14, 18]. Moreover, our work only uses the commonly used traditional PCA and thus does not rely on the time-consuming local PCA as [17] does.

In the next section, we will experimentally show that additional information incorporated in the CFV brings improvement in terms of discriminative power.

3. Experiments

We take the task of material categorization as an example to evaluate our descriptor. As most modern methods have achieved high accuracy [29, 30], say over 96% on the databases such as KTH-TIPS [31], CUReT [32], UIUC [33], and UMD [34], which are shown to contain limited diversity of real-world material appearances, we choose two more challenging datasets:

KTH-TIPS2a dataset (KT2a) [35] contains four samples of eleven different materials, each at nine different scales and twelve different lightings and poses, totalling 4608 images. We follow the setup [35], where three out of the four samples for each material are available for training, and the images of the remaining sample are left for testing.

Flickr Materials Database (FMD) [36] includes ten categories of materials, each of which includes one hundred color images. It was designed for capturing the appearance variations of real-world materials. One characteristic of FMD is the substantial intra-class variation as a result of the diverse selection of instances in each category. It is worth noting that even the human perception can only achieve an accuray of 84.9% on FMD [29]. Thus it is indubitably challenging. According to the protocol by Sharan et al. [36], for each category, we randomly choose 50 images for training and the remaining 50 for testing.

In this section, random experiments are carried out ten times and the average accuracy is reported as the final result on each dataset. Unless specified otherwise we use the one-versus-all SVM with the linear kernel based on LIBSVM [37] with C=1. For the dense SIFT local descriptors and the Fisher vector, we use the implementation of Vlfeat [38]. The gmdistribtuion class in Statistics Toolbox of MATLAB is utilized to optimize the GMM parameters. The symmetric matrix square root is employed to realize $\mathbf{\Sigma}^{-\frac{1}{2}}$ for the CFV, and the only parameter of CFV, α, is set to 0.25 cross all experiments.

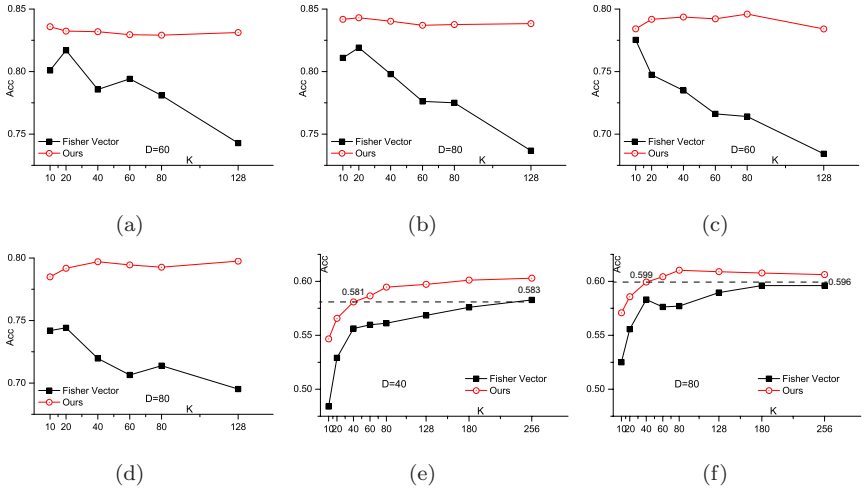

Figure 2: The comparisons between FV and CFV, (a-b) using dense SIFT as local descriptors on KT2a with D=60 and 80 respectively; (c-d) using dense LBP on KT2a with D=60 and 80 respectively; (e-f) using dense SIFT on FMD with D=40 and 80 respectively.

3.1. Completed versus Diagonal

In this subsection, we compare the completed Fisher vector with the FV on various numbers of GMM components. As in [38], all experiments start by doubling the resolution of the input image, and the local descriptors are densely sampled with a factor $\sqrt{2}$ between successive scales, resulting in seven scales from a ratio of 2.0 to 0.25. As the FV, we also use PCA to reduce the dimension of local descriptors for the CFV.

On the KT2a dataset, we employ the following two local descriptors:

Dense SIFT captures a spatial histogram of local gradient orientations. We adopt the default settings of VLFeat [38], i.e. the 128-dimensional dense SIFT with a step of four. The comparisons between our completed Fisher vector with the FV are illustrated by Figure 2, where Figure 2(a) and 2(b) provide the results for $D = 60$ and 80 respectively.

Dense LBP densely calculates the histograms of the 59-bin uniform LBP for a patch of size 16×16 in the RGB color channels separately. Then they are concatenated into a vector representation as the local descriptors. The comparative results are provided by Figure 2(c) and 2(d) for $D = 60$ and 80

respectively.

In line with the results, we have two important observations and related discussions:

1) The completed Fisher vector outperforms the FV under *all* combinations of K, namely the numbers of the GMM components and D, the dimension of the descriptors after PCA regardless of the selection of local descriptors.

2) The KT2a dataset is special as for each material it only contains four samples. Though there are 108 (or 72 for some samples) images under different lighting conditions and scales for each of the four samples, the modes required for fitting these variations are intrinsically limited. It can be observed that for both local descriptors, the optimal K is around 20. The FV framework makes strong assumptions and approximations, which substantially affects the FV by how well the GMM models the data. Thus when K is off the optimal value, the accuracy drops substantially, e.g., by about 10% when K becomes 128. In contrast, as the completed Fisher vector relaxes nearly all assumptions which are adopted for deriving the diagonal FV, the CFV is much more robust to the quality of GMM estimates. As a result, the accuracy of the CFV remains almost the same even when K is varying from 10 to 128.

We turn to the more challenging FMD dataset. Without losing generalization, we use the dense SIFT local descriptors. The comparison results are illustrated by Figure 2(e) and 2(f) for $D = 40$ and 80 respectively. Obviously, more GMM components are required to fit the data since the intra-class variations on FMD are much larger than on KT2a. Nevertheless, the CFV still outperforms the FV for all combinations of D and K, as in the case on KT2a. Moreover, the CFV is also more robust to the changes of K. As a result, by the setting of K is 40, the CFV has achieved comparable or even better accuracy than the FV when $K = 256$.

It is shown that the CFV is insensitive to the number of Gaussians, thus it is able to achieve an acceptable accuracy using much less components than the FV. With a smaller visual vocabulary, the CFV is as efficient as, and even more efficient sometimes than the FV. The reason is that the number of components

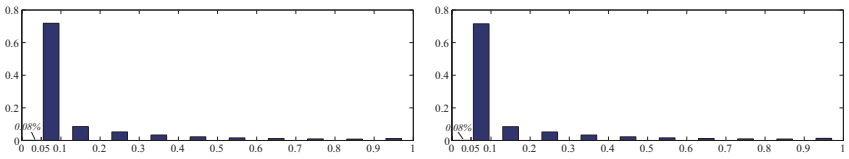

(a) Distributions on training data (b) Distributions on testing data

Figure 3: Distributions of correlation coefficients on FMD dataset

of GMM usually dominates the running time and thus the additional cost for calculating the feature coefficients turns to be trivial.

To further explain the rationality of the CFV, we report the average frequency of the Pearson correlation coefficients based on one particular split into training and test sets on the FMD dataset. More specifically, when the parameter settings of the GMM ($D = 80, K = 256$) are determined, for the dense SIFT local features from each image, we calculated the Pearson correlation coefficients between any two of the feature entries using the soft assignment w.r.t each Gaussian given by Eq. 1 as the probability mass function. The frequency of the absolute values of the correlation coefficients for one image is estimated via a histogram. The average histograms across all components of the GMM from all training and testing images are illustrated in Figure 3(a) and 3(b) respectively. Notice that only a small portion, say 0.08% of the correlation coefficients is below an absolute value of 0.05 (where the two corresponding features are well uncorrelated). For the training data, 91.2% of the correlation coefficients are in the range of $[0.05, 0.5]$, while for the testing data, the portion in the same range becomes 90.8%. The observation shows that the universal PCA fails to eliminate the feature correlations within the Gaussian components. It confirms our doubts about the underlined *local decorrelation* assumption by the Fisher vector framework, and thus suggests the importance of exploring the intra-component feature correlation by the proposed CFV.

3.2. Comparisons with State of the Art

We compare our approach with those state of the art on these two databases. The results of other methods are from corresponding publications. Tables 1

and 2 list the accuracy. It can be found that the CFV with the linear SVM achieves the best accuracy of 84.2% on KT2a and 61.2% on FMD among those methods under consideration. It is also worth noting that on the FMD dataset, the CFV performs as well as the current research hotspot, namely the deep convolutional neural network based approach [39], whereas on KT2a the mulsti-scale CNN representation [40] doesn't show any superiority in performance. Since these two datasets are challenging, the promising results confirm that the feature correlations inside the components of the GMM are informative, and emphasize the superiority of the proposed CFV.

Table 1: Performance on KTH-TIPS2a

Approaches	Models	Accuracy
VZ-MR8 [35]	χ^2 SVM	72.9
VZ-joint [35]	χ^2 SVM	71.0
3-scale LBP^{riu2} [35]	χ^2 SVM	72.1
4-scale LBP^{riu2} [35]	χ^2 SVM	74.5
LHS [41]	*lin.*.SVM	73.0
BOW [30]	*lin.* SVM	74.8
VLAD [30]	*lin.* SVM	76.5
IFV [30]	*lin.* SVM	82.5
CCLBP [42]	χ^2 SVM	77.4
multi-scale CNN [40]	*lin.* SVM	77.4
Ours	*lin.* SVM	**84.2**

4. Conclusions

This paper proposes a completed model of the Fisher vector to remove the limitations brought by the *local decorrelation* assumption in the FV. As an image representation, the CFV encodes the completed statistics by the first two orders of the deviations of the whitened local descriptors w.r.t. the parameters of the GMM. It contains additional information compared with the FV, and

thus leads to enhanced discriminative power. The CFV is more sophisticated than the FV, as the *local decorrelation* assumption is relaxed. In addition, it is insensitive to changes of the number of the GMM components and thus it is able to achieve satisfying accuracy using a relatively small visual vocabulary.

Like the FV, the CFV is efficient. The reduction in the size of vocabulary substantially decreases the running time. Even using the same size of vocabulary, the increase in running time for getting the additional information is moderate, as the inverse square root of the covariance matrix, the largest burden in calculating the CFV given a determined GMM, can be obtained prior to the evaluation procedure. Empirically, we only observe an increase by about 15% in the running time on FMD compared to the FV for $D = 80$ and $K = 256$.

In future, we plan to combine the CFV with the deep convolutional neural network based approaches [43] for further improvement.

5. Acknowledgements

This work was supported by the Academy of Finland and Infotech Oulu.

Table 2: Performance Comparison on FMD

Approaches	Models	Accuracy
Eight-feature [29]	aLDA	39.4
Eight-feature [29]	*int.* SVM	55.6
Eight-feature w.t. mask [29]	aLDA	42.0
Eight-feature w.t. mask [29]	*int.* SVM	57.1
BOW [38]	*lin.* SVM	46.0
VLAD [38]	*lin.* SVM	49.4
FV [38]	*lin.* SVM	59.6
OverFeat. [39]	*lin.* SVM	**61.2**
Ours	*lin.* SVM	**61.2**

References

[1] Y. Huang, Z. Wu, L. Wang, T. Tan, Feature coding in image classification: A comprehensive study, Pattern Analysis and Machine Intelligence, IEEE Transactions on 36 (3) (2014) 493–506.

[2] K. Chatfield, V. Lempitsky, A. Vedaldi, A. Zisserman, The devil is in the details: an evaluation of recent feature encoding methods, in: BMVC, 2011.

[3] T. Ojala, M. Pietikäinen, T. Mäenpää, Multiresolution gray-scale and rotation invariant texture classification with local binary patterns, IEEE Trans. Pattern Anal. Mach. Intell. 24 (7) (2002) 971–987. doi:10.1109/TPAMI.2002.1017623.

[4] D. Lowe, Distinctive image features from scale-invariant keypoints, Int. J. Comput. Vision 60 (2) (2004) 91–110.

[5] K. Mikolajczyk, C. Schmid, A performance evaluation of local descriptors, Pattern Analysis and Machine Intelligence, IEEE Transactions on 27 (10) (2005) 1615–1630.

[6] M. Varma, A. Zisserman, A statistical approach to texture classification from single images, Int. J. Comput. Vision 62 (1-2) (2005) 61–81. doi:10.1007/s11263-005-4635-4.

[7] M. Varma, A. Zisserman, A statistical approach to material classification using image patch exemplars, IEEE Trans. Pattern Anal. Mach. Intell. 31 (11) (2009) 2032–2047. doi:10.1109/TPAMI.2008.182.

[8] X. Tan, B. Triggs, Enhanced local texture feature sets for face recognition under difficult lighting conditions, IEEE Trans. Image Process. 19 (6) (2010) 1635–1650.

[9] H. Bay, A. Ess, T. Tuytelaars, L. Van Gool, Speeded-up robust features (SURF), Computer vision and image understanding 110 (3) (2008) 346–359.

[10] S. ul Hussain, B. Triggs, Visual recognition using local quantized patterns, in: Computer Vision–ECCV 2012, Springer, 2012, pp. 716–729.

[11] J. Sánchez, F. Perronnin, T. Mensink, J. Verbeek, Image classification with the fisher vector: Theory and practice, International journal of computer vision 105 (3) (2013) 222–245.

[12] O. Boiman, E. Shechtman, M. Irani, In defense of nearest-neighbor based image classification, in: Computer Vision and Pattern Recognition, 2008. CVPR 2008. IEEE Conference on, IEEE, 2008, pp. 1–8.

[13] F. Perronnin, C. Dance, Fisher kernels on visual vocabularies for image categorization, in: Computer Vision and Pattern Recognition, 2007. CVPR'07. IEEE Conference on, IEEE, 2007, pp. 1–8.

[14] H. Jégou, M. Douze, C. Schmid, P. Pérez, Aggregating local descriptors into a compact image representation, in: Computer Vision and Pattern Recognition (CVPR), 2010 IEEE Conference on, IEEE, 2010, pp. 3304–3311.

[15] X. Zhou, K. Yu, T. Zhang, T. S. Huang, Image classification using supervector coding of local image descriptors, in: Computer Vision–ECCV 2010, Springer, 2010, pp. 141–154.

[16] D. Picard, P.-H. Gosselin, Improving image similarity with vectors of locally aggregated tensors, in: Image Processing (ICIP), 2011 18th IEEE International Conference on, IEEE, 2011, pp. 669–672.

[17] J. Delhumeau, P.-H. Gosselin, H. Jégou, P. Pérez, Revisiting the vlad image representation, in: Proceedings of the 21st ACM international conference on Multimedia, ACM, 2013, pp. 653–656.

[18] F. Perronnin, J. Sánchez, T. Mensink, Improving the fisher kernel for large-scale image classification, in: Computer Vision–ECCV 2010, Springer, 2010, pp. 143–156.

[19] F. Perronnin, Y. Liu, J. Sánchez, H. Poirier, Large-scale image retrieval with compressed fisher vectors, in: Computer Vision and Pattern Recognition (CVPR), 2010 IEEE Conference on, IEEE, 2010, pp. 3384–3391.

[20] K. Simonyan, O. M. Parkhi, A. Vedaldi, A. Zisserman, Fisher vector faces in the wild, in: Proc. BMVC, Vol. 1, 2013, p. 7.

[21] J. Krapac, J. Verbeek, F. Jurie, Modeling spatial layout with fisher vectors for image categorization, in: Computer Vision (ICCV), 2011 IEEE International Conference on, IEEE, 2011, pp. 1487–1494.

[22] X. Peng, C. Zou, Y. Qiao, Q. Peng, Action recognition with stacked fisher vectors, in: Computer Vision–ECCV 2014, Springer, 2014, pp. 581–595.

[23] X. Peng, L. Wang, X. Wang, Y. Qiao, Bag of visual words and fusion methods for action recognition: Comprehensive study and good practice, CoRR abs/1405.4506.

[24] P.-H. Gosselin, N. Murray, H. Jégou, F. Perronnin, Inria+ xerox@ fgcomp: Boosting the fisher vector for fine-grained classification, `http://hal.inria.fr/docs/00/93/91/04/PDF/RR-8431.pdf/` (2013).

[25] R. Duda, P. Hart, D. Stork, Pattern classification (2nd Eds.), John Wiley & Sons, 2000.

[26] H. Jégou, O. Chum, Negative evidences and co-occurences in image retrieval: The benefit of PCA and whitening, in: Computer Vision–ECCV 2012, Springer, 2012, pp. 774–787.

[27] O. Chapelle, P. Haffner, V. N. Vapnik, Support vector machines for histogram-based image classification, Neural Networks, IEEE Transactions on 10 (5) (1999) 1055–1064.

[28] A. Vedaldi, A. Zisserman, Efficient additive kernels via explicit feature maps, Pattern Analysis and Machine Intelligence, IEEE Transactions on 34 (3) (2012) 480–492.

[29] L. Sharan, C. Liu, R. Rosenholtz, E. H. Adelson, Recognizing materials using perceptually inspired features, International journal of computer vision 103 (3) (2013) 348–371.

[30] M. Cimpoi, S. Maji, I. Kokkinos, S. Mohamed, A. Vedaldi, Describing textures in the wild, in: Proc. of IEEE Conference on Computer Vision and Pattern Recognition, 2014.

[31] B. Caputo, E. Hayman, M. Fritz, J.-O. Eklundh, Classifying materials in the real world, Image Vision Comput. 28 (1) (2010) 150–163.

[32] K. Dana, B. van Ginneken, S. Nayar, J. Koenderink, Reflectance and texture of real-world surfaces, ACM Trans. Graph. 18 (1) (1999) 1–34.

[33] S. Lazebnik, C. Schmid, J. Ponce, A sparse texture representation using local affine regions, Pattern Analysis and Machine Intelligence, IEEE Transactions on 27 (8) (2005) 1265–1278.

[34] Y. Xu, H. Ji, C. Fermüller, Viewpoint invariant texture description using fractal analysis, International Journal of Computer Vision 83 (1) (2009) 85–100.

[35] B. Caputo, E. Hayman, P. Mallikarjuna, Class-specific material categorisation, in: Proc. IEEE International Conference on Computer Vision, 2005, pp. 1597–1604.

[36] L. Sharan, R. Rosenholtz, E. Adelson, Material perception: What can you see in a brief glance?, Journal of Vision 9 (8) (2009) 784–784.

[37] C.-C. Chang, C.-J. Lin, Libsvm: a library for support vector machines, ACM Transactions on Intelligent Systems and Technology (TIST) 2 (3) (2011) 27.

[38] A. Vedaldi, B. Fulkerson, VLFeat: An open and portable library of computer vision algorithms, `http://www.vlfeat.org/` (2008).

[39] P. Sermanet, D. Eigen, X. Zhang, M. Mathieu, R. Fergus, Y. LeCun, Overfeat: Integrated recognition, localization and detection using convolutional networks, arXiv preprint arXiv:1312.6229.

[40] W. Li, Learning multi-scale representations for material classification, in: Pattern Recognition, Springer, 2014, pp. 757–764.

[41] G. Sharma, S. ul Hussain, F. Jurie, Local higher-order statistics (lhs) for texture categorization and facial analysis, in: Proc. of ECCV, Springer, 2012.

[42] X. Qi, Y. Qiao, C.-G. Li, J. Guo, Exploring cross-channel texture correlation for color texture classification, in: Proc. of the BMVC, 2013.

[43] K. Chatfield, K. Simonyan, A. Vedaldi, A. Zisserman, Return of the devil in the details: Delving deep into convolutional nets, in: British Machine Vision Conference, 2014. arXiv:1405.3531.

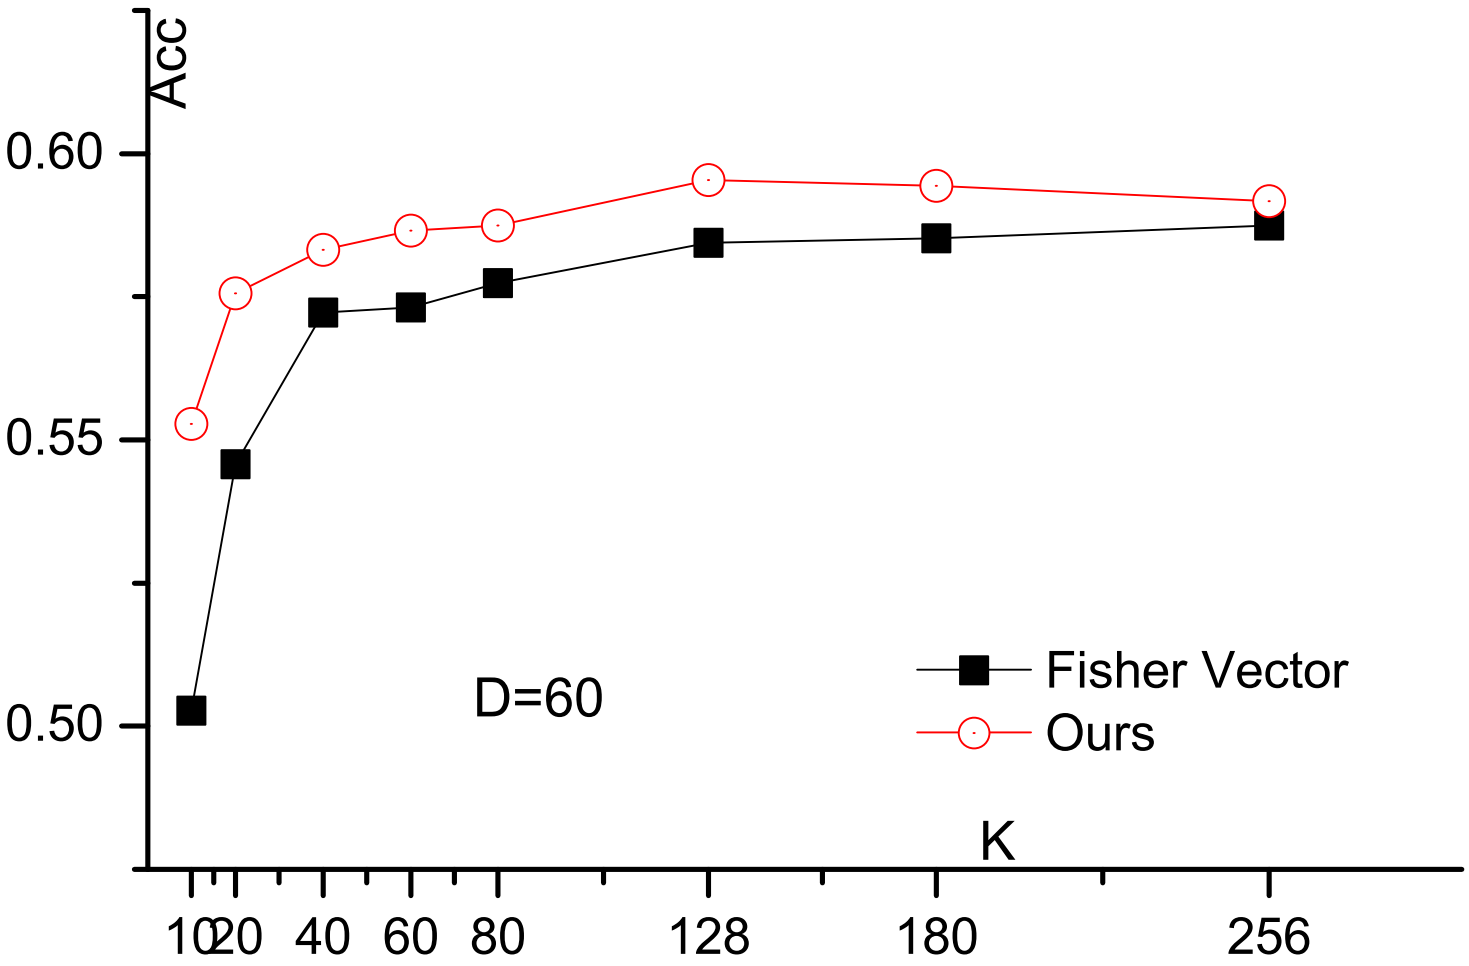
Acc
0.60
0.55
0.50
D=60
Fisher Vector
Ours
K
10
20
40
60
80
128
180
256